THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 59, No. 6, July–August 1980
Printed in U.S.A.

Trunk Demand Servicing in the Presence of Measurement Uncertainty

By C. R. SZELAG

(Manuscript received October 24, 1979)

Demand servicing attempts to correct existing overload conditions in the trunk network in a timely and cost-effective manner. The demand servicing procedures presented in this paper achieve this goal by taking actions that have minimal effect on the existing network configuration, account for the statistical nature of the traffic measurements, and follow the network design objectives. In comparison to existing methods, these procedures should reduce unnecessary servicing activity and improve the consistency between demand servicing and planned servicing.

I. INTRODUCTION

Trunk network administration is composed of two major operations, planned servicing and demand servicing. Together, they determine the quantities and locations of interoffice trunks required to maintain a reasonable balance between network service and cost. Planned servicing provides circuits, using forecasts of demand made on a yearly basis. Demand servicing uses recent traffic measurements to detect and correct existing service problems. Because the Bell System trunk network represents a multi-billion dollar investment in equipment, personnel, and operation support systems, it is important that these functions be cost-effective and that they be performed efficiently. This paper describes recently developed procedures for demand servicing the trunk network.

Demand servicing compensates for the effects of forecast errors on the planned network. Its main function is to detect the need for trunk group augments when service problems develop, and to take corrective action in a timely and cost-effective manner. In general, demand servicing is not concerned with disconnecting trunks in excess of current demand, since the removal of excess trunks is a part of the planned servicing function, which makes such decisions on the basis of

the year-to-year trunk forecast. Demand servicing represents a minimal adjustment necessary to restore service, rather than a complete reconfiguration of the network.

The demand-servicing procedure described in this paper is designed for implementation as a new feature in the Trunk Servicing System (TSS), a standard software tool used by Bell System operating companies to monitor the performance of the trunk network. It represents an integral part of an overall plan for network administration.[1]

Section II of this paper reviews the fundamentals of trunk network design and introduces some of the problems to be considered. Section III discusses the demand-servicing policy objectives that motivated the approach taken in the design of the servicing procedure described in Sections IV and V.

II. BACKGROUND AND MOTIVATION

This section reviews the fundamentals of trunk network design and administration and introduces the key issues involved in this study.

2.1 Statistical nature of demand servicing

The trunk servicer relies on traffic measurements to monitor network service and decide on corrective action when overload conditions are detected. To illustrate the statistical nature of this process, consider the simple, two-node network shown in Fig. 1a. The trunk group joining end offices A and B provides the only route for calls originating at A and destined for B. Such an "only-route" trunk group is sized according to an average blocking criterion, where blocking refers to the fraction of calls which arrive, but fail to find an idle circuit to their destination. The service objective for an only-route trunk group is 0.01 average blocking in the time-consistent busy hour of the busy season.

Using a forecast of traffic anticipated in the next busy season, planned servicing provides sufficient circuits to carry the expected traffic at the 1-percent blocking objective. To monitor trunk group service, the statistic $\bar{B}_n = 1/n \sum_{i=1}^{n} B_i$ is computed for the busy hour, where B_i represents the fraction of calls blocked in the ith day of an n-

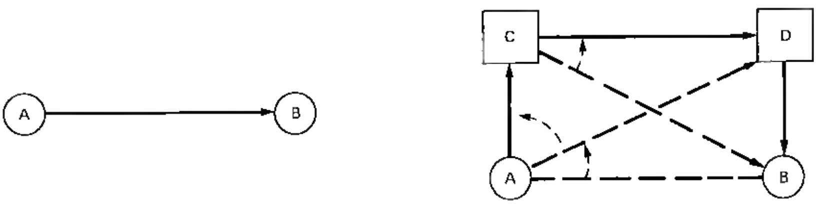

(a) (b)

Fig. 1—(a) Only-route trunk group. (b) Alternate-routing network.

day study period.* Because of the finite number of days in the sample and the finite (1-hour) measurement interval each day, the average blocking will deviate from the 0.01 objective even when the group is correctly sized.[1]

Because of the importance of maintaining good service and the reluctance to disconnect trunks in excess of current demand, overreaction to these statistically volatile estimates can lead to unnecessary trunk group augments and higher-than-necessary levels of reserve capacity. Thus, it is important that the trunk servicer be provided with statistically reliable procedures for monitoring network service.

2.2 Network considerations

The demand servicing problem is considerably more difficult in a complex network environment. To understand the issues involved, we first review the fundamental principles of trunk network design.

For reasons of economics, most end-office pairs are connected by a complex network that allows for the alternate routing of calls, rather than the simple only-route trunk group described above. Consider, for example, the network shown in Fig. 1b. Calls originating at A and destined for B are first offered to the primary high-usage (PH) group AB. Failing to find an idle circuit, the call is then alternate-routed to the intermediate high-usage (IH) group AD. If an idle circuit is available, the call is switched at D to its destination; otherwise, it is again alternate-routed to the final group AC. If an idle circuit is available, the call then proceeds to search for a path to B. If it is not possible to establish a path from source to destination, an "all trunks busy" signal (usually a recorded message or a fast busy signal) is given to the calling party, who must retry at a later time.

Note that two types of trunk groups are shown here. Final groups (represented by solid lines) are sized for the same 0.01 average blocking objective described above for only-route groups. Also, note that as long as blocking is low on the final groups, calls originating at office A have a high probability of completion.

High-usage trunk groups (represented by dashed lines) are usually designed for a much higher rate of overflow. Such a group is sized to balance the incremental load-carrying costs between its direct and alternate routes.

For demand servicing in an alternate routing network, the following questions must be answered:

(i) Are the traffic measurements on a trunk group consistent with the group's service objective, or are additional circuits required?

---

* A 20-day study period consisting of four consecutive business weeks is recommended, although smaller samples frequently are obtained.

(*ii*) If corrective action is to be taken, which groups in the network should be augmented?

(*iii*) How many trunks should be added?

The development of a trunk demand servicing policy that effectively addresses these issues is described in the remainder of this paper.

III. DEMAND SERVICING POLICY OBJECTIVES

The goal of demand servicing can be stated simply: Restore objective network service in a timely and cost-effective manner. This goal can be accomplished by a servicing policy designed to achieve the objectives discussed in this section.

3.1 Account for the variability of traffic estimates

We have already seen (Section 2.1) how performance statistics based on traffic measurements are used to monitor network service. To avoid unnecessary trunk group augments and their associated costs, allowance must be made for the statistical variability of these estimates. Conversely, it is also important to detect poor service when such conditions exist. Thus, demand servicing rules must account for this variability to maintain a reasonable balance between cost and service.

3.2 Minimize network activity

Demand servicing action is required only when planned servicing fails to provide acceptable network service. It is important that this restoration of service be carried out quickly. Also, the need for the equipment required to provide the additional circuits probably has not been anticipated by the trunk forecast, the basic input to the equipment planning process. Therefore, these circuits may be more difficult to obtain. For these reasons, demand servicing should result in a minimal amount of network activity, both in the number of groups affected and in the number of additional circuits required to restore service.

Since the blocking experienced by traffic originating from a network cluster* is ultimately determined by the blocking on the final trunk group, it is only necessary to demand-service a network cluster when the blocking on the final group is significantly greater than the 0.01 service objective. Also, a high-usage group in an overloaded cluster should only be considered if it is contributing excess overflow traffic to the overloaded final.

To summarize, we attempt to minimize the number of trunk groups affected by demand-servicing action by focusing on the overloaded

---

* A network cluster consists of all high-usage trunk groups which originate at a common switching office and overflow to a common final trunk group. See Fig 2.

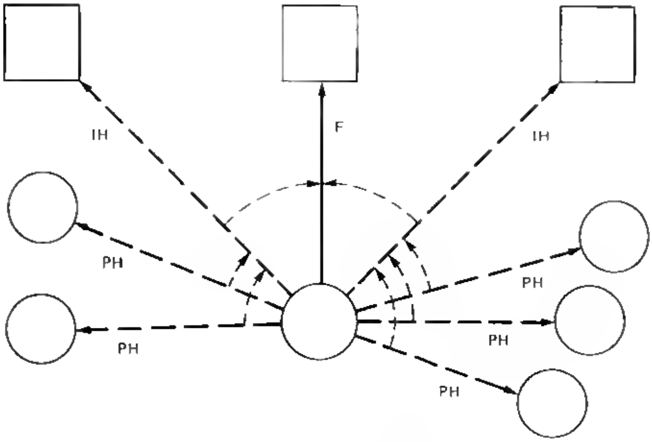

Fig. 2—Example network cluster.

network cluster final and those subtending high-usage groups contributing to the high-blocking conditions.

3.3 Service low in the network

When a significant service problem has been detected on the final group of a network cluster, the group's blocking can be reduced by either

(*i*) Adding trunks to the final.

(*ii*) Reducing the load offered to the final by augmenting subtending high-usage groups, thus reducing their overflow.

(*iii*) Some combination of the above.

Although it may appear that augmenting the final is the easiest way to restore service, there are several reasons for avoiding this type of action.

We first consider the immediate effects of routing extra traffic via the alternate-route final rather than on the direct high-usage groups. As an overflow call rises in the network hierarchy, its path tends to become longer and to involve more switching. Thus, it is usually more expensive to carry and complete a multi-switched call rather than a direct call. Also, traffic carried on the lower levels of the network has less chance to interact with (e.g., block) other traffic in the system. Therefore, when demand servicing is required, trunks should be added to the groups that are currently underprovided, based on economic engineering considerations.

Now consider the long-term effects of servicing exclusively high in the trunking hierarchy, e.g., the final. The result of this practice is that the final is actually overprovided by engineering standards, while high-usage groups are undersized. Thus, the yearly planned servicing proc-

ess would likely stimulate considerable rearrangements, representing both capital costs for equipment and significant administrative and labor expenses, in order to restore the network to a minimum-cost configuration.

These considerations motivate a servicing policy that attempts to resolve the service problem at the lowest level of the trunking hierarchy consistent with the excess traffic causing the overload.

3.4 Implement in TSS

The demand-servicing algorithm developed in this study was designed for implementation in the current TSS system. This imposed constraints on both the type of information available and the way the algorithm would be implemented (i.e., batch rather than interactive processing).

IV. SERVICING THRESHOLDS

As we have already seen, an understanding of the statistical variability of the traffic measurements that drive the demand servicing process is essential to the formulation of an effective servicing policy. In this section, we use previous work that quantifies this variability to develop statistical tests for detecting overloaded groups.

4.1 Thresholds of acceptable blocking for grade of service trunk groups

For trunk groups sized for an average blocking objective,* the key statistic used to monitor service is the observed blocking $\bar{B}_n$ over an n-day study period. Simulation studies and analysis by Neal[1] have determined the distribution of this statistic. The distribution depends on the number of trunks in the group, the number of days in the average, and the characteristics of the call arrival process—the mean load $\bar{a}$, its within-hour variation or peakedness z, and its level of day-to-day variation ϕ. Figure 3 shows the average blocking distributions of a correctly engineered grade-of-service trunk group for both 10- and 20-day study periods.

Given the distribution of $\bar{B}_n$, it is possible to develop a simple statistical test to decide whether the measured average blocking is significantly greater than the 0.01 objective. The test uses a threshold or upper bound B_u of acceptable observed blocking. When $\bar{B}_n$ exceeds B_u, we decide to take corrective action; otherwise, the measured blocking is acceptable and no action is taken. In this way, the upper bound B_u defines the acceptable deviation from the service objective.

The choice of the threshold B_u is designed to achieve a reasonable

---

* Such trunk groups are called "grade-of-service" groups. These include the only-route and final groups previously discussed.

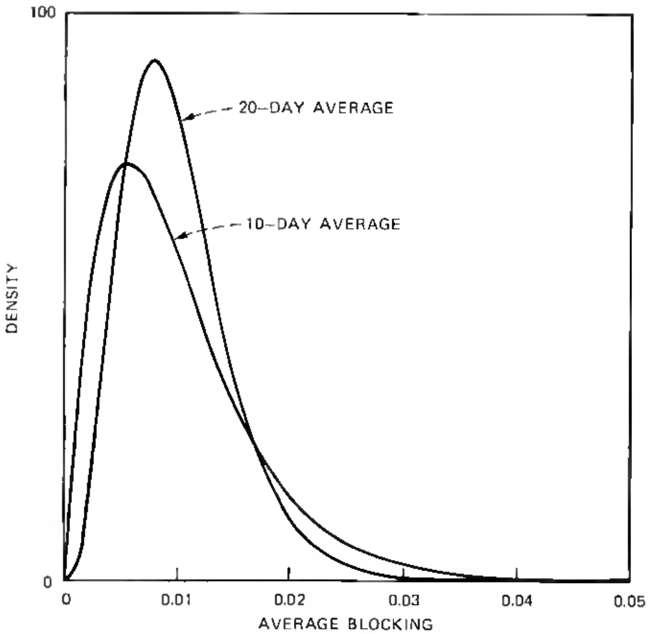

Fig. 3—Variability of measured blocking.

balance between two types of servicing errors. False alarms, or Type I errors, occur when the measured blocking $\bar{B}_n$ exceeds the threshold even though the group is correctly sized. Misses, or Type II errors, occur when the group is overloaded, but the measurement falls below the threshold and the problem is not detected. As we can see in Fig. 4, raising the threshold decreases the false-alarm probability but increases the miss probability. Although we have not quantified the costs associated with these errors, relative differences between the two can be used to establish performance criteria.

Each false alarm results in the unnecessary expenditure of both capital and labor on an emergency basis. Because of the lead time required to react to a service problem (typically, several weeks to a few months after it is detected), it is possible that the busy season may have passed by the time circuits can be added. Thus, the trunk servicer should be reasonably certain that a significant problem exists before issuing an order for additional circuits. This means that the false-alarm probability should be fairly small, i.e., a few percent.

Conversely, a miss occurs when the underlying true mean blocking exceeds the objective, but the realized blocking (i.e., the blocking experienced by the customer) falls below the threshold B_u. However, if a real service problem is missed in one study period, it may still be

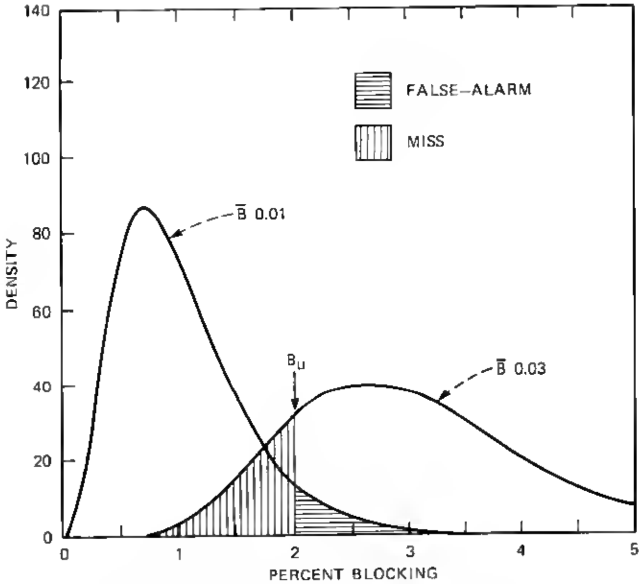

Fig. 4—False-alarm and miss probabilities.

detected in subsequent study periods. For these reasons, misses are relatively less serious than false alarms.

Guided by these considerations, the following performance criteria were selected for determining the blocking thresholds:

(*i*) A false alarm probability ≤ 0.025.
(*ii*) A miss probability ≤ 0.10 when the group is overloaded to a 0.05 blocking level.

These criteria apply to study-period averages based on at least 15 days of valid measurements, which are classified as A-quality data in the TSS system.

If B_1 is the 0.975-quantile of the observed blocking distribution for a correctly sized group, then any $B_u \geq B_1$ satisfies condition (*i*). If B_2 is the 0.10-quantile of the blocking distribution for a group with 5-percent mean blocking, then any $B_u \leq B_2$ satisfies condition (*ii*). Thus, if $B_1 \leq B_2$, it is possible to satisfy both criteria.

For all but the smallest trunk groups (i.e., fewer than eight trunks), $B_1 \leq B_2$ and both criteria may be satisfied. The thresholds selected for these groups are shown in Table I. They depend only on trunk group size and satisfy both criteria over the range of traffic conditions (values of z and ϕ) encountered in practice. For the smallest trunk groups, the thresholds satisfy condition (*i*) only.

For study periods with fewer than 15 days of data, it was generally

not possible to satisfy both criteria. To avoid the unnecessary servicing of trunk groups based on the more volatile estimates derived from these smaller samples, the thresholds of acceptable blocking are larger and satisfy condition (*i*) only. In this way, the trunk servicer, who is responsible for providing good service, is also encouraged to base his servicing decisions on high-quality data. Costly decisions which cannot be justified from the data are avoided.

4.2 Thresholds for high-usage groups

This section develops trunk estimate thresholds for high-usage groups which complement those developed in the previous section for grade-of-service groups.

A high-usage trunk group sized to balance the incremental load-carrying costs of the direct and alternate routes has its load carried on the last trunk equal to its economic CCS (ECCS) value.[3] Based on this criterion, a group's offered load and load characteristics estimated from study period traffic measurements are used to compute the number of trunks required, $\hat{c}_R$, to achieve this economic balance. As is the case for grade-of-service groups sized for an average blocking objective, the statistical variability of these trunk estimates must also be taken into account.

This variability has been quantified by Hill,[2] who developed a normal approximation to the distribution of $\hat{c}_R$. Using those results, it is possible to derive thresholds of acceptable high-usage trunk estimates that satisfy specified criteria. Specifically, given the standard deviation $\sigma_{\hat{c}_R}$ of $\hat{c}_R$ (see the appendix) and an objective false-alarm probability α, the threshold for considering servicing action on a high-usage group with c_I trunks in service is $c_I + k_\alpha \cdot \sigma_{\hat{c}_R}$, where k_α depends only on α and is commonly tabulated.[4] Thus, it remains to specify only the false-alarm probability α.

While it is important to prevent false alarms, it is also important to detect overloaded groups. As we showed in Section 4.1, these are

Table I—Thresholds of acceptable measured blocking for grade-of-service trunk groups (c = number of trunks in group, N = number of days in average)

	B_u			
c	$15 \le N \le 20$	$11 \le N \le 14$	$7 \le N \le 10$	$3 \le N \le 6$
2	0.070	0.080	0.090	0.140
3	0.050	0.060	0.070	0.090
4	0.050	0.060	0.070	0.080
5–6	0.040	0.050	0.060	0.070
$6 \le c \le 336$	0.030	0.035	0.040	0.060
$336 \le c \le 504$	0.025	0.030	0.035	0.055
$504 \le c$	0.020	0.025	0.030	0.050

conflicting objectives. Consequently, the mutual effects of the false-alarm and detection probabilities must be considered.

Recall that one of the demand servicing objectives is to identify the source of overload conditions and resolve the problem as low in the network as is justified. Therefore, when overload conditions are detected on the cluster final due to undersized high-usage groups lower in the network, it is important that these groups contributing to the problem be identified and serviced. Thus, the choice of false-alarm probabilities for high-usage groups should provide adequate probabilities of detection.

Detection probabilities are a function of the degree of overload. Recall that the thresholds of acceptable service for finals were designed to detect groups having 0.05 blocking with at least 90-percent probability. Typically, 0.05 blocking occurs when the offered load is approximately 25 percent larger than the design load. Thus, 0.05 blocking on the final can occur if the overflow from all the subtending groups increases by 25 percent; accordingly, the high-usage detection procedure should identify groups with 25-percent excess overflow.

Similarly, an increased load to the final of roughly 15 percent results in 0.03 blocking. Hence, if the detection process identifies half the groups with 25-percent overloads and they are augmented, the new load to the final will have an expected blocking of less than 0.03, and the final-group blocking will most likely be acceptable.

To select the upper limit of acceptable estimates, loads which produced approximately 25-percent excess overflow over a wide range of trunk-group size, peakedness, level of day-to-day variation, and ECCS were generated. For each trunk-group size, tbe detection probabilities for given false-alarm levels were averaged over the other parameters. The results indicated that, as trunk-group size increases, so does the detection probability, but that for a false-alarm probability $\alpha = 0.10$, even up to high-usage groups with 48 trunks, the average detection probability is less than 0.40. For high-usage groups with 24 or more trunks, a false-alarm probability of 0.2 provides the desired probability of detection. Since the small groups are less likely to be causing a problem on the final because of their smaller overflow, a false-alarm probability of 0.2 satisfied our objectives and, hence, will be used to specify the servicing thresholds for high-usage groups.

To summarize, a high-usage trunk group is considered for servicing when $\hat{c}_R > c_I + T$, where $T = k_\alpha \sigma_{\hat{c}_R}$, $\alpha = 0.20$, and $k_\alpha = 0.8416$. The appendix gives an approximation to $\sigma_{\hat{c}_R}$ as a function of group size, data quality, and level of day-to-day load variation.

4.3 Significant trunk requirement for high-usage groups

Suppose that c_I trunks are in service and the true (expected) requirement is $\bar{c}_R$. According to our model, $\Pr\{\bar{c}_R \geq \hat{c}_R - T\} \approx 0.8$, where $\hat{c}_R$

is the estimated trunk requirement. This means that, with high probability, at least $\hat{c}_R - T$ trunks are truly required. The quantity $c_R^S = \hat{c}_R - T$ represents a conservative estimate of the true requirement $\bar{c}_r$ and will be called the significant trunk requirement. Its use is discussed in Section 5.3.

4.4 Summary of servicing thresholds

Upper limits (thresholds) of observed average blocking and trunks-required estimates were developed above for grade-of-service and high-usage trunk groups, respectively. These thresholds are used to identify overload conditions on individual trunk groups while allowing for statistical variation in the estimates.

V. DEMAND SERVICING ALGORITHM

We now present a demand-servicing algorithm for network clusters. Our procedure consists of four parts:

(*i*) Deciding whether servicing is required.

(*ii*) Identifying groups in the cluster which are contributing to the problem (these groups are called candidates for servicing).

(*iii*) Determining augments for candidates.

(*iv*) Deciding which candidates should be serviced.

This sequence of decisions is executed in a cyclic order until the problem has been resolved. Each cycle results in either the decision that no action is required or that one of the candidate groups be serviced with its recommended augment. Details are given below.

5.1 Identifying clusters for servicing

A network cluster is selected for servicing when its observed grade of service, as measured by the average blocking on the final group during the study period, is significantly greater than the 0.01 objective. This decision is made using the statistical test described in Section 4.1. If the measured blocking exceeds the upper limit B_u, corrective action should be taken.

5.2 Selection of candidates for servicing

Using the trunk estimate thresholds developed in Section 4.2, individual high-usage groups in the cluster can be tested for overload conditions. Some, but perhaps not all, these groups will be selected as candidates for additional circuits by the procedure described below.

As previously indicated, we recommended that the network be serviced at the lowest effective level, servicing only those high-usage groups that are contributing excess overflow to the final. These groups are chosen by starting at the highest level of the trunking hierarchy, the final, and proceeding downward into the cluster as follows. At each level of the hierarchy, only those undertrunked groups subtending

undertrunked groups at the next higher level are selected as candidates for servicing. Thus, we start at the (overloaded) final group and proceed downward until we find either a group with acceptable service or a primary high-usage group. This means that an overloaded primary high-usage (PH) group subtending an intermediate high-usage (IH) group with acceptable performance is not selected as a servicing candidate.

5.3 Recommended number of trunks for candidates

After a high-usage or final trunk group is selected as a candidate for servicing, another decision on the actual number of circuits to add must be made. This should depend not only on the estimate of trunks required, $\hat{c}_R$, but also on the following considerations:

(*i*) The availability of additional facilities for the group.

(*ii*) Forecasts of requirements.

(*iii*) The servicing of subtending groups.

However, specific information on the availability and costs of additional circuits is not, in general, available in a mechanized fashion to a trunk servicing system such as TSS. Therefore, the number of trunks planned or forecast for the imminent busy season, c_F, may be used as a rough indication of the availability of facilities. However, we do not want to use c_F as a strict upper bound on the number of trunks to be considered. Demand servicing should allow for unanticipated growth in the network while maintaining a near-minimum cost configuration.

Conversely, suppose that the estimated (current) trunk requirement is significantly greater than the number of trunks in service, but less than that forecast for the current period ($c_I + T < \hat{c}_R < c_F$). Then one might be tempted to augment the group to its full forecasted requirement, c_F, to anticipate future demand and to avoid both a service problem and the possibility of an additional augment in the near future. However, the forecast c_F available to a system like TSS is typically based on data at least a year old and may not accurately reflect demand in the near future, recent traffic transfers, and other relevant considerations. Hence, the practice of automatically servicing a group up to its forecast requirement c_F is of questionable merit, especially in the context of demand servicing.

Another factor to consider in the specification of a recommended trunk group size for groups that receive overflow traffic is the servicing of subtending groups. Adding trunks to subtending groups will decrease the offered load to the group under consideration; hence, fewer trunks will be required. If this reduced requirement, $\tilde{\hat{c}}_R$, still exceeds the acceptable level $c_I + T$, the recommended number of trunks, c'_{Rec}, is determined from this revised estimate by the method described below.

If $\hat{c}'_R$ does not exceed the threshold, the group should not be serviced.

Motivated by the considerations above, we recommend that an overloaded, high-usage trunk group (selected as a servicing candidate by the above rule) be sized to c_{Rec}, where

$$c_{Rec} = \begin{cases} \hat{c}_R & \text{if} \quad \hat{c}_R \le c_F \\ \max(c_F, c_R^S) & \text{if} \quad \hat{c}_R > c_F, \end{cases}$$

and $c_R^S = \hat{c}_R - T$ is the statistically significant trunk requirement.

In other words, when the estimated trunk requirement exceeds the forecast, only the statistically significant number c_R^S are recommended. However, if $c_R^S < c_F < \hat{c}_R$, then the forecast c_F should be implemented.

We see in Section 5.4 that the selection of high-usage groups for servicing will be based on the magnitude of the reduction in overflow reaching the final when c_{Rec} trunks are implemented. Since this reduction in overflow is related to the magnitude of the trunk group augment, consideration of the smaller, conservative estimate c_R^S of the number of trunks required when $\hat{c}_R > c_F$ will favor the selection of those groups for which additional facilities should be available.

The recommended augment for final trunk groups is given in the servicing algorithm in Section 5.5.

5.4 Selecting groups

The decision to service the cluster is made when the measured average blocking on the final, $\bar{B}_n$, exceeds the threshold B_u. Using the "top down" approach described in Section 5.2, we identify those high-usage groups that appear to be contributing to the overload condition. For each of these candidates for additional trunks, an (initial) augment, $c_{Rec} - c_I$, is specified.

Recall that a major objective of our demand servicing policy is to resolve the service problem at its source. This is accomplished by starting low in the network and working up into the cluster until the blocking on the final group is at an acceptable level. Specifically, we start by selecting an overloaded primary high-usage group, if one has been identified as a candidate, and augmenting it to its recommended level, c_{Rec}. Since we would like to service as few groups as possible, we select, at each level of the trunking hierarchy, the candidate whose servicing will contribute the maximum reduction in load offered to the final. After this group is augmented, we determine whether additional action is required (Section 5.1), reexamine the candidacy of all groups affected by the augment (Section 5.2), and select the next group, if necessary.

This procedure is described in detail below.

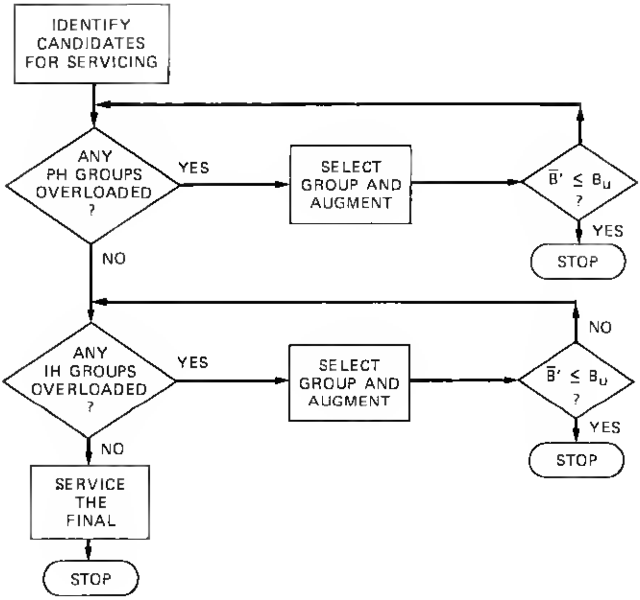

Fig. 5—Cluster servicing algorithm.

5.5 Servicing algorithm (see Fig. 5)*

1. Service overloaded PH groups (if no PH groups were initially selected as candidates for servicing, or if none remain as candidates, go to 2):
   (*a*) For each (remaining) candidate PH group, compute Δl, the load that will be removed from the final if the number of trunks in the group is increased from c_I to c_{Rec}.
   (*b*) Choose the PH group with the largest Δl value computed in (*a*) as the group to be serviced.
   (*c*) Recompute the average blocking on the final, $\bar{B}'$. If $\bar{B}' \le B_u$, STOP. Otherwise, recompute trunk requirements on all groups receiving overflow from the group just chosen, and reevaluate the candidacy of all high-usage groups affected by the servicing action. (These include all high-usage groups that receive overflow from the group just serviced and all candidate high-usage groups in the same subcluster as the group just serviced.)

* The algorithm described in this section assumes that only one level of IH groups is present in the network cluster trunking hierarchy. For more general networks, with IH groups overflowing to other IH groups, the algorithm can be generalized by stratifying the (overloaded) IH groups according to terminating office class. Step 2 is executed for each set of IH groups terminating at the same office class, starting with the lowest class and proceeding upward in the network.

2. Service IH groups (if no IH groups were initially selected as candidates or if none remain after 1, go to 3):
   (a) For each (remaining) candidate IH group, compute Δl as defined in 1a.
   (b) Choose for servicing the IH group with the largest Δl value computed above. If none remain, go to 3.
   (c) Recompute average blocking on the final, $\bar{B}'$. If $\bar{B}' \le B_u$, STOP. Otherwise, go to b.
3. Service the final: Using the (revised) estimate of load offered to the final, compute the number of trunks required for 0.01 average blocking. STOP.

The algorithm terminates when either the servicing of subtending high-usage groups has lowered the final's blocking to an acceptable level (less than B_u) or the final itself has been serviced.

VI. SUMMARY

This paper has described a demand servicing procedure designed for implementation in the Trunk Servicing System (TSS). Our approach consists of three parts: identifying network clusters which require servicing, locating specific groups within the cluster to be serviced, and determining how these groups should be augmented.

The demand servicing procedure described here can be used to achieve the objective network performance in a cost-effective manner, allowing for such considerations as the statistical properties of the traffic measurements, the effect of servicing at different levels of the trunking hierarchy, and the availability of facilities. Following these procedures should reduce costly unnecessary network servicing activities and ensure that servicing action is consistent with the network design objectives.

VII. ACKNOWLEDGMENT

The author would like to acknowledge the contributions of S. R. Neal, whose work on the statistical variability of traffic estimates[1]

Table II—Curve fits to standard deviations of trunk estimates
$\sigma(c) = b_0(l) + b_1(l)c$

Level of Load Variation	Coefficient	
l	$b_0(l)$	$b_1(l)$
Low	0.578	0.019
Medium	0.459	0.035
High	0.305	0.053

played a fundamental role in the development of the demand servicing procedure described in this paper.

APPENDIX

Standard Deviation of Trunk Estimates for High-Usage Groups

Using the method described in Ref. 2, the standard deviation $\sigma_{\hat{c}_H}$ of high-usage trunk estimates was computed under the assumption that the trunk group is engineered to provide a specified economic load on the last trunk (ECCS). In these studies, we examined trunk-group sizes from 3 to 240, ECCS values from 6 to 24, traffic peakedness from 1.0 to 4.0, and level of day-to-day load variation from low to high. It was assumed that hourly measurements of usage, peg count, and overflow (UPCO) were taken over a 20-day study period.

The results indicated that, for each level of day-to-day load variation, the standard deviation is linearly related to the number of trunks in service but is rather insensitive to peakedness and ECCS. A least-squares fit to the data for each level of day-to-day variation is accurate within one trunk for trunk-group sizes up to 240 trunks. Thus, the standard deviations can be considered to be functions of two variables, trunk-group size c, and level of day-to-day variation l.

The linear approximations $\sigma(c)$ to the standard deviations, computed for each level of day-to-day variation, are given by

$$\sigma(c) = b_0(l) + b_1(l)c,$$

where $b_0(l)$ and $b_1(l)$ were computed by a least-squares regression for l = low, medium, or high and are given in Table II.

Because $\sigma_{\hat{c}_H}$ is computed for 20-day samples, the factor $\sqrt{20/n_d}$, where n_d is the number of days in the trunk study, is used to correct for non-20-day study periods.[2]

In general,

$$\sigma_{\hat{c}_H} = \sqrt{20/n_d}\,[b_0(l) + b_1(l)c].$$

REFERENCES

1. S. R. Neal, "Blocking Distributions for Trunk Network Administration," B. S. T. J., this issue, pp. 829–844.
2. D. W. Hill, "The Effects of Traffic Load Variation on Measurement Accuracy," B. S. T. J., *56,* No. 4 (April 1977), pp. 561–573.
3. C. J. Truitt, "Traffic Engineering Techniques for Determining Trunk Requirements in Alternate Routing Trunk Networks," B. S. T. J., *31,* No. 2 (March 1954), pp. 277–302.
4. M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions,* Applied Mathematics Series 55, Washington, D.C., National Bureau of Standards, June 1964.